

# Temporal Attention for Cross-View Sequential Image Localization


Dong Yuan<sup>1</sup>, Frederic Maire<sup>1</sup>, Feras Dayoub<sup>2</sup>

**Abstract**—This paper introduces a novel approach to enhancing cross-view localization, focusing on the fine-grained, sequential localization of street-view images within a single known satellite image patch, a significant departure from traditional one-to-one image retrieval methods. By expanding to sequential image fine-grained localization, our model, equipped with a novel Temporal Attention Module (TAM), leverages contextual information to significantly improve sequential image localization accuracy. Our method shows substantial reductions in both mean and median localization errors on the Cross-View Image Sequence (CVIS) dataset, outperforming current state-of-the-art single-image localization techniques. Additionally, by adapting the KITTI-CVL dataset into sequential image sets, we not only offer a more realistic dataset for future research but also demonstrate our model’s robust generalization capabilities across varying times and areas, evidenced by a 75.3% reduction in mean distance error in cross-view sequential image localization.


## I. INTRODUCTION

Cross-view image geo-localization endeavours to enhance the precision of robot localization in outdoor environments that may have erratic GPS signals. Most previous works [1], [2], [3], [4], [5], [6], [7], [8], [9] formulate the cross-view geo-localization problem as an image retrieval task, and solve it by metric learning. When provided with a current street-view image captured by the robot, they retrieve the corresponding satellite image patch from a reference dataset, assigning the center GPS coordinate of the retrieved satellite patch as the location of the street-view image. Despite achieving a commendable level of accuracy in image retrieval, these results rely on a one-to-one setting where one street-view image is exactly located at the center of one satellite image from the reference dataset. In real-world scenarios, it is impractical to guarantee that the current street-view image precisely aligns with the center of a specific satellite image within the reference dataset. Consequently, employing the center GPS coordinate of the retrieved satellite image as the current location invariably leads to estimation errors.

Departing from the idea of image retrieval, recent works [10], [11], [12] focus on the cross-view fine-grained localization task, which involves predicting pixel coordinates corresponding to the current street view on a given satellite image patch (see Fig.1). Zhu *et al.* [12] introduced the VIGOR dataset for research purposes and proposed a regression-based approach to predict the offset of the street-view location relative to the center point of the satellite



Cross-view fine-grained localization



Cross-view sequential image localization

Fig. 1. Comparison of two cross-view localization tasks. **Top:** Cross-view fine-grained localization primarily emphasizes discrete location predictions, where street-view images are widely spaced apart and exhibit discontinuous distribution along the route or area of interest. **Bottom:** Cross-view sequential image localization is centered around predicting the location in a satellite image of each street-view image belonging to a sequence, aligning more closely with practical localization applications.

image. Xia *et al.* [11] employed the encoding of satellite images and street-view images into local and global descriptors, respectively. Using correlation operations between the global descriptor and each local descriptor, a dense location distribution map is derived. Yuan *et al.* [10] utilized cross-attention mechanisms to establish correspondences between two views. They then integrated classification and regression prediction methods, resulting in state-of-the-art performance on the cross-view fine-grained localization task.

We argue that the aforementioned cross-view fine-grained localization studies predominantly concentrate on the localization of individual street-view images. However, in practical robot localization applications, continuous streams of ground-view images are typically received from the robot’s cameras. Consequently, extending the cross-view fine-grained localization task to sequential ground-view image localization is more practical and more acceptable in real-world scenarios.

In this paper, we refer to the novel task as *cross-view*

<sup>1</sup>QUT Centre for Robotics, Queensland University of Technology, Australia {yuand2, f.maire}@qut.edu.au

<sup>2</sup>Australian Institute for Machine Learning (AIML), University of Adelaide, Australia feras.dayoub@adelaide.edu.au

*sequential image localization*, involving the prediction of sequential ground-view image location within a known satellite image patch, as illustrated in Fig.1. To leverage temporal information effectively, we adopt the state-of-the-art method on the cross-view fine-grained task [10] as the foundational framework and augment it by adding a hidden state to retain information from the preceding time step. Subsequently, we introduce a novel temporal attention module designed to fuse information from both the hidden state and the current state, thereby generating updated features for the prediction of the current street-view location. We evaluate our method on a large-scale cross-view image sequence dataset (CVIS) [13], which contains thousands of sequential ground-view images with corresponding satellite images. We also partition the KITTI-CVL dataset [14] into numerous ground-view image sequences, facilitating the assessment of our model’s generalization across different times and unknown areas. Our main contributions can be summarized as follows:

- We introduce a novel temporal attention mechanism that enhances the model’s ability to utilize historical context from sequential ground-view images for more accurate localization.
- We enhance the KITTI-CVL dataset by segmenting it into realistic short image sequences, creating a refined dataset for evaluating cross-view sequential image localization. This effort, coupled with our use of the CVIS dataset, underpins our experiments, showing our method’s improved accuracy and robustness.

To foster future research on cross-view sequential image localization, we make our code available at: <https://github.com/UQ-DongYuan/CVSeqLocation>

## II. RELATED WORK

### A. Cross-view Image Retrieval

Cross-view image retrieval methodologies commonly entail the extraction of high-dimensional feature vectors, referred to as image embeddings, from ground and aerial view images. The determination of matching between images from different views is achieved through the computation of the distance between their respective image embeddings. Hence, the pivotal factor in enhancing retrieval accuracy lies in the extraction of high-quality features from ground and aerial view images. CVM-Net [1] employed NetVLAD [15] layers to extract view-invariant features to improve performance. Shi *et al.* [5] introduced a feature transport module designed to align features extracted from both view images. Bridging the domain gap between ground and aerial views can also lead to an enhancement in retrieval accuracy. Regmi *et al.* [3] employed a conditional Generative Adversarial Networks (cGAN) [16] to synthesize aerial images from ground-view images, subsequently extracting features from the synthesised images. SAFA [2], on the other hand, leveraged prior geometric knowledge to convert aerial images into ground-view by utilizing polar transformation, and proposed a spatial-aware module to enhance the spatial correlation of extracted features. Toker *et al.* [17] combined SAFA [2] with

a cGAN to achieve promising results. Building upon the success of Transformer [18] in image recognition [19] and object detection [20], [6] and [9] leveraged Transformer as the backbone for feature extraction from both ground-view and aerial view images. Through the incorporation of position embedding and self-attention mechanisms [18], the extracted features are enriched with spatial information, consequently improving the accuracy of retrieval.

While the aforementioned methods have demonstrated high retrieval accuracy, their accuracy in location estimation is limited by the assumption of center alignment and the setting of one-to-one matching.

### B. Cross-view Fine-grained Localization

To address the limitations inherent in cross-view image retrieval, Zhu *et al.* [12] established the VIGOR dataset, wherein each street-view image is covered by multiple satellite images and positioned non-centrally within the satellite image. They explored the application of regression methods to predict the offset of street-view images relative to the center of satellite images. Building upon the work of SAFA [2], Xia *et al.* [11] extracted multiple local descriptors as well as a global descriptor from satellite and street-view images, respectively. By conducting correlation operations between these descriptors, a similarity score map is generated. This map is then upsampled to obtain a location probabilistic distribution map of the same dimensions as the satellite image, facilitating the prediction of pixel locations of street-view images within satellite images. Yuan *et al.* [10] denoted this task as cross-view fine-grained localization, introducing the cross-attention mechanism to compute the cross-attention score between extracted features from two views. Leveraging this score, they obtained fusion features that integrate correspondences between the two views. Finally, employing a combination of classification and regression akin to object detection tasks, they further reduce the location prediction errors.

Our proposed method extends the foundational framework [10] by introducing a temporal attention module to handle temporal information, enabling the prediction of pixel-wise locations for each street-view image within an image sequence.

### C. Cross-view Image Retrieval Using Image Sequence

Previous cross-view image retrieval methods focused on matching a single street-view image with satellite images in a reference dataset. [14] and [13] extended the one-to-one image retrieval work by using continuous street-view images to match a satellite image. [14] proposed CVLNet to match a continuous frame of street-view images to a large satellite image for camera localization. The CVLNet transforms and projects sequential ground-view images into an overhead view and utilizes photo consistency among the projected images to form a global representation. To evaluate their model, they collected corresponding satellite images from Google Maps for the KITTI dataset [21] to construct the KITTI-CVL benchmark dataset. Similarly, [13] employed a

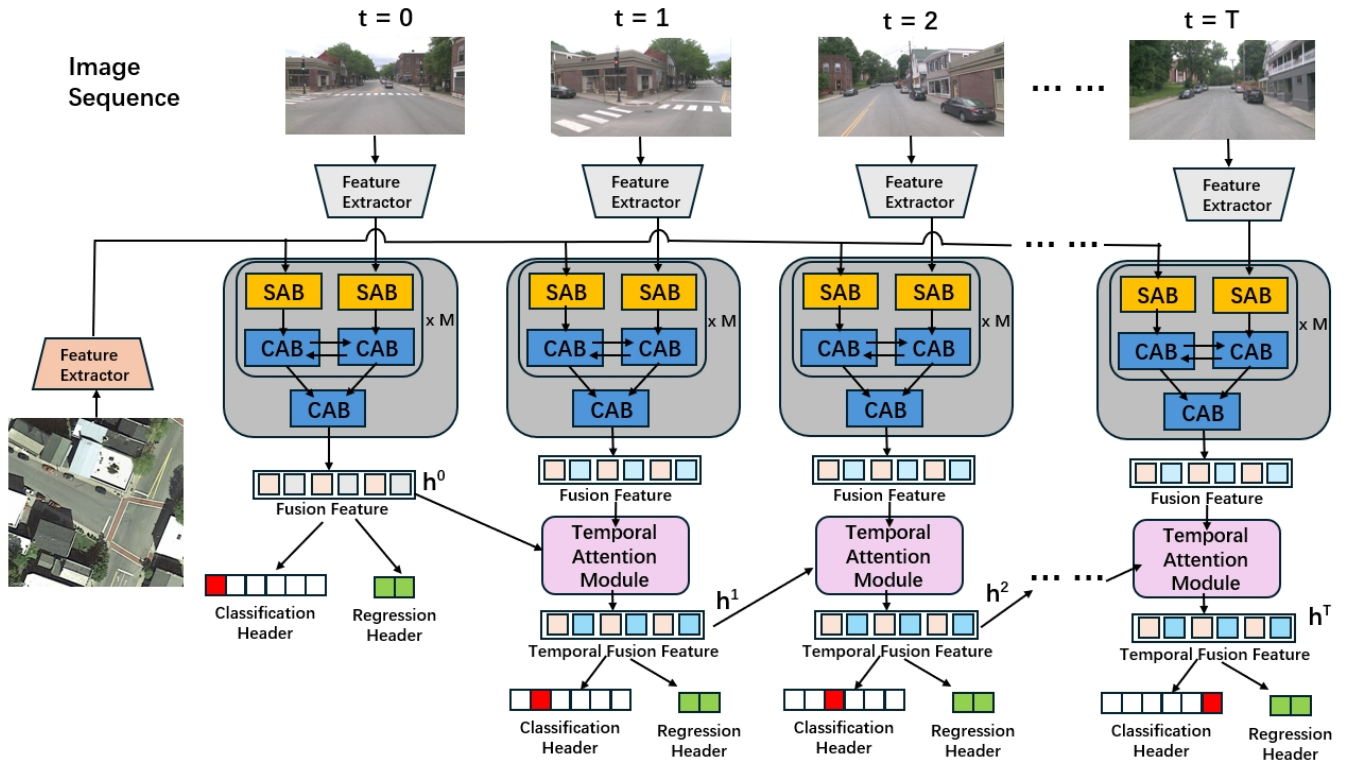


Fig. 2. Overview of the proposed architecture. The feature extractors extract feature maps independently from two different views. These feature maps are then fed into self-attention blocks (SAB) and cross-attention blocks (CAB) to generate fused features. The fusion features are further integrated with the hidden state  $h^{t-1}$  from the previous time step within the temporal attention module. Finally, the obtained features are utilized for the prediction of locations.

Transformer structure to fuse multiple consecutive ground-view images into a global descriptor for matching with a satellite image. Additionally, they also utilized street-view images captured by on-board cameras and satellite images sourced from Google Maps to build a large-scale cross-view image sequence geo-localization (CVIS) dataset. While these two studies are rooted in image retrieval and aim to match a street-view image sequence with a satellite image, our proposed novel task involves predicting the pixel coordinates of sequential street-view images in a satellite image. Theoretically, our new task is more useful for robot localization deployed in the real world. We evaluate the performance of our approach using the aforementioned two datasets, which will be detailed further in the Experiment Settings Section IV.

### III. METHODOLOGY

In this section, we present the proposed method for cross-view sequential image localization. We introduce a method that achieves state-of-the-art performance on cross-view fine-grained localization [10] as our baseline framework in Section III-A. Section III-B introduces the proposed *Hidden State*, which retains information from the preceding time step. In Section III-C, we present a novel Temporal Attention Module (TAM), which can integrate information from both the previous and current time steps. An overview of the proposed framework is presented in Fig.2. Section III-D introduces the training objectives.

#### A. Baseline Framework

Given that the proposed method in [10] yielded state-of-the-art outcomes in the cross-view fine-grained localization task, we adopted this method as our baseline framework. This framework primarily comprises local feature extractors, the cross-view feature fusion network, and prediction headers, which will be elaborated upon in the subsequent section.

1) *Local Feature Extractor*: Two modified ResNet50 [22] neural networks are employed to extract features from the satellite image and ground-view image, respectively. Upon passing the ground-view image through the feature extractor, a feature map  $F_g \in \mathbb{R}^{C \times H_g \times W_g}$  is derived, where  $H_g$  and  $W_g$  are  $1/8$  of the height and width of the original input ground-view image, and  $C$  is the channel dimension. Similarly, the satellite image undergoes feature extraction, resulting in a feature map  $F_s \in \mathbb{R}^{C \times N \times N}$ , where  $N$  represents the dimensions of the extracted satellite image feature map. Subsequently, the resulting feature maps from two views are flattened, and a  $1 \times 1$  convolution layer is utilized to reduce the channel dimension  $C$  to  $D$ . The obtained features serve as input for the subsequent stage.

2) *Cross-View Feature Fusion Network*: The cross-view feature fusion network primarily consists of self-attention blocks and cross-attention blocks, utilized to enhance and fuse image features, respectively. Features extracted from the ground-view and satellite images,  $F_g \in \mathbb{R}^{D \times H_g \times W_g}$  and  $F_s \in \mathbb{R}^{D \times N^2}$ , are fed into a self-attention block (SAB)

followed by a cross-attention block (CAB), respectively. In the cross-attention block, the features from one view (either ground-view or aerial view) provide Key and Value, while the features from the other view provide Query to conduct multi-head cross-attention operations. This process (an SAB followed by a CAB) repeats  $M$  times and concludes with a CAB to derive the ultimate fusion feature  $F_f \in \mathbb{R}^{D \times N^2}$ .

3) *Prediction Headers*: Two prediction headers are constructed on the fusion features: a classification header and a regression header. Both prediction headers consist of three-layer perceptrons. The classification header produces an  $N^2 \times 1$  vector, tasked with predicting the specific cell of the  $N \times N$  grid where the street-view image is located. Meanwhile, the regression header is designed to predict the offsets associated with the grid selected by the classification header.

### B. Hidden State

Recurrent Neural Networks (RNNs) [23] are a class of neural networks specifically designed to process sequential data by maintaining an internal state or memory. This memory or internal state enables RNNs to capture temporal dependencies in sequential data, making them successful in tasks such as speech recognition [24] and text generation [25]. In the Cross-view sequential image localization task, we argue that two consecutive street-view images share overlapping semantic information. Consequently, leveraging information from the preceding time step is advantageous for accurately predicting the location at the current time step. Inspired by the idea of RNNs, we create a hidden state  $h$  to retain information from the preceding time step. At time  $t$ , the satellite image  $I_s^t$  and ground-view image  $I_g^t$  are processed through local feature extractors and a cross-view feature fusion network (detailed in Section III-A), yielding the fusion feature  $F_f^t$ . Subsequently, the fusion feature  $F_f^t$  and the hidden state  $h^{t-1}$  from the previous time step are fed into our proposed Temporal Attention Module (to be elaborated upon in the subsequent section), resulting in the temporal information integrated feature  $\hat{F}_f^t$ . This derived feature is utilized to connect the prediction headers for location prediction. Concurrently, the derived feature  $\hat{F}_f^t$  is stored as the hidden state  $h^t$  at the current time step and will contribute to the temporal attention process at the next time step. At the first time step  $t = 0$ , unlike conventional RNNs, we do not initialize a hidden state with a zero vector. Instead, we directly store the fusion feature  $F_f^0$  as the hidden state  $h^0$  and utilize  $F_f^0$  to predict the location of the first ground-view image. The aforementioned procedure can be described as follows:

$$h^t = \begin{cases} F_f^0 & t = 0 \\ \text{TAM}(F_f^t, h^{t-1}) & t > 0 \end{cases} \quad (1)$$

### C. Temporal Attention Module (TAM)

The proposed Temporal Attention Module aims to integrate temporal information from the previous time step into

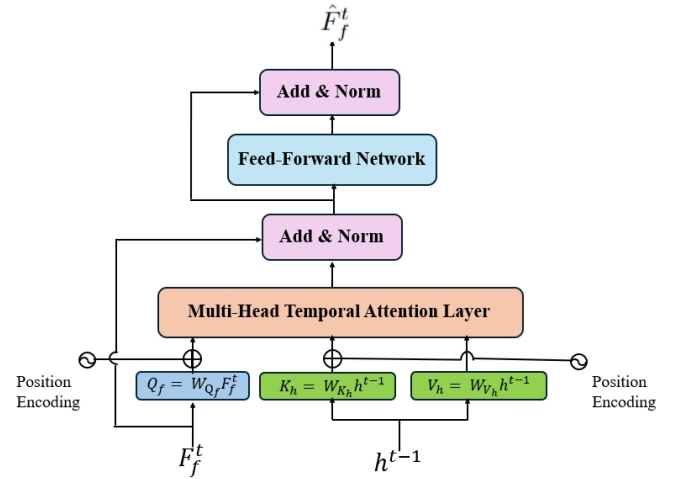


Fig. 3. Illustration of Temporal Attention Module (TAM). The fusion feature  $F_f^t$  undergoes projection to generate the Query  $Q_f$ , while the hidden state  $h^{t-1}$  from the previous time step is projected to form the Key  $K_h$  and Value  $V_h$ . The Position Encoding is applied to both  $Q_f$  and  $K_h$  before they are passed into the multi-head attention layers.

the current prediction process, shown in Fig.3. We achieve this by using the attention mechanism [18], where Q (query), K (key) and V (value) are fundamental components involved in the calculation of the attention mechanism. The objective is to capture the correlation between the current fusion feature  $F_f^t$  and the hidden state  $h^{t-1}$  through attention score computation. This correlation (the attention score) allows us to extract relevant information from the hidden state  $h^{t-1}$ , which is subsequently integrated with the fusion feature  $F_f^t$  to derive a new feature  $\hat{F}_f^t$ . Hence, in this temporal attention process, the current fusion feature  $F_f^t$  is projected to Q, while the hidden state  $h^{t-1}$  is projected to both K and V. This projection process can be presented as:

$$\begin{aligned} Q_f &= W_{Q_f} F_f^t \\ K_h &= W_{K_h} h^{t-1} \\ V_h &= W_{V_h} h^{t-1} \end{aligned} \quad (2)$$

where  $W_{Q_f}$ ,  $W_{K_h}$  and  $W_{V_h}$  are matrices to be learned. Following the approach in [20], before multi-head attention layers, sinusoidal positional encodings  $E_{pos}$  are added to  $Q_f$  and  $K_h$  respectively to maintain the positional information. Based on the attention score, relevant information  $\tilde{F}_{tem}$  extracted from hidden state  $h^{t-1}$  can be obtained by as follows:

$$\tilde{F}_{tem} = \text{softmax}((Q_f + E_{pos})^T (K_h + E_{pos})) V_h \quad (3)$$

The relevant information  $\tilde{F}_{tem}$  is added to the original fusion feature  $F_f^t$  and undergoes processing through a feed-forward network (FFN), comprising two linear projection layers. Finally, the output feature  $\hat{F}_f^t$  is obtained for position prediction and concurrently stored as the current hidden state  $h^t$ .

#### D. Training Objective

Consistent with the approach outlined in [10], we adopt the standard categorical cross-entropy loss function for training the classification header mentioned in III-A.3, which is defined as:

$$\mathcal{L}_{cls} = - \sum_i^C y_i \log(\text{softmax}(p_i)) \quad (4)$$

where  $y_i$  represents the ground-truth label indicating the index of grid cells in an  $N \times N$  grid and  $p_i$  denotes the output of the classification header. The mean square error (MSE) is employed for training the offset regression header, which can be formulated as follows:

$$\mathcal{L}_{mse} = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2 \quad (5)$$

where  $y_i$  denotes the ground-truth offset values and  $\hat{y}_i$  denotes the prediction offset values. In contrast to the approach proposed in [10], our method receives a continuous sequence of street-view images as input. Consequently, we compute the average of the sequence loss as the objective loss for each iteration, which can be formulated as follows:

$$\mathcal{L}_{seq} = \frac{1}{T} \sum_{t=1}^T \lambda_{cls} \mathcal{L}_{cls}^t + \lambda_{mse} \mathcal{L}_{mse}^t \quad (6)$$

where  $T$  denotes the number of street-view images in a sequence.  $\lambda_{cls}$  and  $\lambda_{mse}$  are the corresponding weight coefficients.

### IV. EXPERIMENTAL SETTINGS

#### A. Datasets

We conduct experiments on two datasets, namely CVIS [13] and KITTI-CVL [14]. We choose the method proposed in [10] as our baseline and compare the performance on the cross-view sequential image localization task.

1) *CVIS Dataset*: The CVIS dataset consists of street-view images captured by a front-camera mounted on a vehicle in urban and suburban areas of the state of Vermont, USA. Each street-view image in the dataset has a resolution of  $1920 \times 1080$  pixels and is accompanied by accurate GPS coordinates and heading direction information. Within a street-view image sequence, the distance between consecutive images is approximately 8 meters, while the distance between the first and last images does not exceed 50 meters. This constraint is implemented to ensure that a street-view image sequence can be covered within a satellite image. In total, the dataset comprises 38863 street-view image sequences, with each sequence containing an average of 7 images. Additionally, to represent more real-world scenarios, around 70% of street-view images are collected from suburban areas and 30% from urban areas.

Satellite images matching the number of street-view sequences are acquired using Google Maps. The GPS coordinates of the street-view image located in the middle of each sequence are selected and then randomly adjusted



Fig. 4. A sequence sample segmented from the KITTI-CVL dataset. Each red marker denotes the location of a street-view image. The red arrow indicates the direction of travel of the vehicle.

by a 5-meter offset to determine the coordinates of the corresponding satellite image. Each aerial image is captured at a zoom level of 20, yielding a resolution of  $640 \times 640$  pixels. The ground resolution of satellite images is approximately 0.114 meters per pixel.

We split the CVIS dataset into training, validation and testing sets, comprising 24871, 6217, and 7773 aerial-ground sequence pairs, respectively.

2) *KITTI-CVL dataset*: The KITTI dataset holds considerable significance in autonomous driving and computer vision research [21]. Shi *et al.* [14] undertook the task of gathering satellite images that correspond to street-view images from Google Earth, thereby enriching the KITTI-CVL Dataset to investigate cross-view geo-localization challenges. Each satellite image in this dataset maintains a resolution of  $1280 \times 1280$  pixels, with an associated ground resolution of approximately 20 centimetres per pixel.

In the original KITTI dataset, consecutive street-view images are highly similar with a small distance between them. To better suit practical applications of sequential image localization, we adopt a methodology akin to the CVIS dataset [13] by segmenting each complete video into multiple continuous street-view sequences. Specifically, we extract a frame every 8 meters from the entirety of the original video, generating a lengthy sequence denoted as  $S = s_0, s_1, \dots, s_x$ . We commence by computing the distance between  $s_0$  and  $s_1$ . If this distance falls below 50 meters, we proceed to calculate the distance between  $s_0$  and  $s_2$ , iteratively checking until the distance between  $s_0$  and  $s_x$  exceeds 50 meters. Subsequently,  $[s_0, s_x]$  constitutes a sequence. The next sequence starts from the midpoint of  $[s_0, s_x]$ . Upon deriving all sequences, those containing fewer than 6 street-view images are discarded, culminating in a final set of 1077 street-view image sequences. The satellite image corresponding to the midpoint location of each sequence is collected, yielding 1077 pairs of satellite images and corresponding street-view image sequences. A sample pair of ground-aerial images after being segmented from the KITTI-CVL dataset is shown in Fig.4. Similarly, all aerial-ground sequence pairs

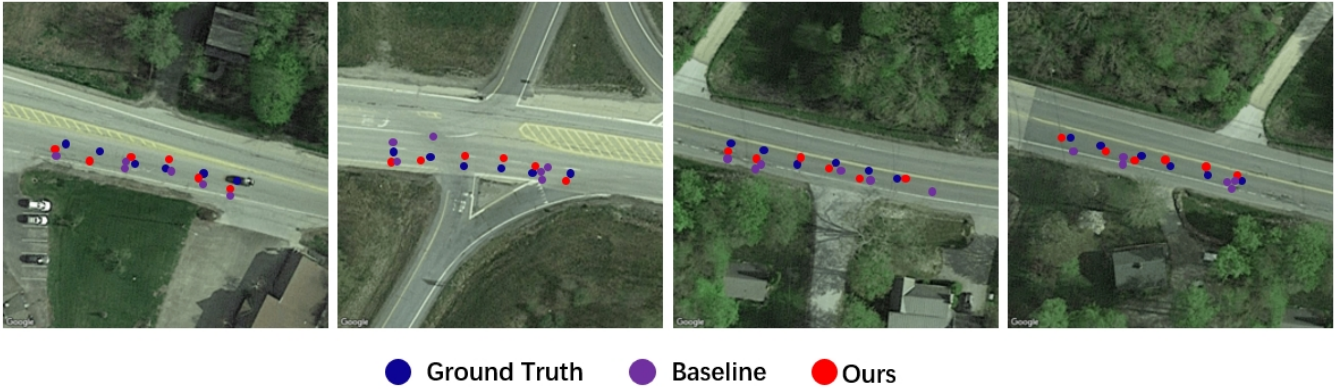


Fig. 5. Qualitative localization result visualization on the CVIS [13] dataset. We use blue, purple and red points to denote the predicted consecutive street-view locations of ground truth, the baseline method and our proposed method. The figure illustrates that the baseline method tends to cluster multiple predicted locations in the same area, whereas our approach demonstrates a closer alignment with the actual ground truth locations.

are subdivided into 862 pairs for training, 107 pairs for validation, and 108 pairs for testing.

### B. Implementation Details

For the CVIS dataset, both satellite and street-view images are resized to dimensions of  $256 \times 256$  and  $270 \times 480$ , respectively, before being fed into the model as input. Regarding the KITTI-CVL dataset, the satellite images undergo an initial center cropping to  $640 \times 640$  before being resized to  $256 \times 256$ . Similarly, the street-view images are resized to  $145 \times 480$  to prepare them for model input. We follow the settings in [10] to construct the baseline framework, where the channel dimension  $C = 1024$  of the feature maps obtained from the feature extractors, and will be reduced to  $D = 256$  through a  $1 \times 1$  convolutional layer. The dimension  $N$  of the satellite image feature map and the classification grid  $N \times N$  are both set to 32. The SAB followed by the CAB process will be repeated  $M = 4$  times. We initialize the parameters of the local feature extractor using pre-trained weights from ResNet50 [22], trained on the ImageNet dataset [26]. The remaining parameters are initialized using Xavier initialization [27]. For model optimization, we employ the AdamW optimizer [28]. Specifically, we set the learning rate for the feature extractors to  $10^{-5}$ , while the learning rate for other parameters is set to  $10^{-4}$ .

1) *Training Process*: Firstly, we initiate training for the baseline method [10]. This involves randomly selecting a street-view image from each sequence, pairing it with the corresponding satellite image covering the sequence, and feeding the pair into the model. The batch size is configured to 4. During inference, the entire sequence is inputted, and the location of each street-view image will be predicted. Subsequently, we proceed to train our proposed new method. Here, we utilize the pre-trained feature extractors from the baseline method, keeping them frozen, to extract feature maps. The obtained feature maps will be fed into the subsequent training stage for training parameters of other components. Each input comprises a street-view image sequence paired with its corresponding satellite image, with a batch size of 1. Each sequence contains 6 consecutive street-view

TABLE I  
COMPARISON BETWEEN OUR METHOD WITH SINGLE-IMAGE FINE-GRAINED LOCALIZATION METHODS ON THE CVIS DATASET. THE RESULTS ARE SHOWN ON THE VALIDATION AND TEST DATASETS. BEST RESULTS ARE IN BOLD.

	Validation		Test	
	mean	median	mean	median
Localization error(m)				
CVML [11]	15.99	14.72	16.37	15.13
Baseline [10]	12.31	11.84	12.57	12.15
Ours	<b>4.30</b>	<b>1.89</b>	<b>3.29</b>	<b>1.74</b>

images for training and testing. Regarding the KITTI-CVL dataset, the limited number of sequences obtained through segmentation raises concerns about overfitting during direct training. Therefore, we opt to fine-tune the model trained on the CVIS dataset using the KITTI-CVL data.

2) *Evaluation Metrics*: In our evaluation process, we predict the locations of all street-view images within each sequence and subsequently compute the average distance error for the entire sequence, denoted as the sequence error. Upon collecting all sequence errors within the dataset, we derive both the mean and median values of all sequence errors. These two measures serve as evaluative metrics to assess the performance of the methods.

## V. EXPERIMENTS RESULTS

### A. Results on CVIS Dataset

The results from the comparison of cross-view sequential image localization on the CVIS dataset are summarized in Table I. Apart from the baseline method, another cross-view fine-grained localization method [11] mentioned in Section II-C is trained and tested on the CVIS dataset for comparison with our proposed method. As depicted in Table I, the performance of these two methods considering the location of a single street view is notably poor in predicting the location of continuous street-view images. This poor performance can be attributed to the similarity among street-view images in a sequence. Models that only consider a

TABLE II

COMPARISON BETWEEN OUR METHOD WITH THE BASELINE METHOD ON THE KITTI-CVL DATASET. THE RESULTS ARE SHOWN ON THE VALIDATION AND TEST DATASETS. BEST RESULTS ARE IN BOLD.

Localization error(m)	Validation		Test	
	mean	median	mean	median
Baseline [10]	13.56	13.17	13.91	13.47
Ours	11.37	10.53	12.45	11.91
Ours + Fine-tuning	<b>3.00</b>	<b>1.70</b>	<b>3.07</b>	<b>2.10</b>

single image struggle to establish a robust correspondence between similar street views and different locations, leading to larger prediction distance errors.

The visualization results (see Fig.5) further illustrate that some street-view images, which are relatively similar, are erroneously predicted to be in the same area by these models. In contrast, the newly proposed method incorporates a temporal attention module that compares and integrates information from the previous time step with that of the current time step. This mechanism enhances the accuracy of the model in predicting sequential locations. Compared to the baseline method, our proposed method achieves a significant improvement in both mean distance error and median distance error on the test dataset, with enhancements of 73.8% and 85.7%, respectively.

### B. Generalization in Unknown Area across Times

The generalization ability of our method is assessed on unknown regions and different dates using the KITTI-CVL dataset. The street-view images in the KITTI dataset utilized for training and testing were collected across various times and areas. The experimental results are presented in Table II.

In Table II, the first two rows display the outcomes of the baseline method and our new method directly trained on the KITTI-CVL dataset. It is evident that our new method outperforms the baseline method. Remarkably, when fine-tuning the model trained on the CVIS dataset with the KITTI-CVL dataset, there is a significant improvement in prediction accuracy. This indicates that our method exhibits strong generalization ability across different location areas and dates. Moreover, our model achieves commendable performance across diverse data domains through fine-tuning with a small amount of data, highlighting its generalization capability across different domains.

## VI. CONCLUSION AND FUTURE WORK

In this work, we expand upon the cross-view fine-grained localization problem to address the cross-view sequential image localization challenge, which aligns more closely with real-world applications. To tackle this task, we introduce a novel temporal attention module designed to integrate information from both the preceding and current time steps, thereby enhancing the accuracy of predicting the position of street-view images at the current step. Furthermore, we partition the KITTI-CVL dataset into shorter street-view image

sequences, facilitating the evaluation of cross-view sequence image localization. Through extensive experimentation, our proposed method demonstrates notable effectiveness in predicting sequential street-view image locations. Additionally, it exhibits superior generalization capabilities in unknown areas across different times. Future work will incorporate motion prior information to enhance the accuracy of cross-view sequential image localization in large-scale areas.

## REFERENCES

- [1] S. Hu, M. Feng, R. M. Nguyen, and G. H. Lee, "Cvm-net: Cross-view matching network for image-based ground-to-aerial geo-localization," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7258–7267.
- [2] Y. Shi, L. Liu, X. Yu, and H. Li, "Spatial-aware feature aggregation for image based cross-view geo-localization," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [3] K. Regmi and M. Shah, "Bridging the domain gap for ground-to-aerial image matching," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 470–479.
- [4] L. Liu and H. Li, "Lending orientation to neural networks for cross-view geo-localization," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 5624–5633.
- [5] Y. Shi, X. Yu, L. Liu, T. Zhang, and H. Li, "Optimal feature transport for cross-view image geo-localization," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 07, 2020, pp. 11 990–11 997.
- [6] S. Zhu, M. Shah, and C. Chen, "Transgeo: Transformer is all you need for cross-view image geo-localization," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 1162–1171.
- [7] Y. Shi, X. Yu, D. Campbell, and H. Li, "Where am i looking at? joint location and orientation estimation by cross-view matching," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 4064–4072.
- [8] S. Cai, Y. Guo, S. Khan, J. Hu, and G. Wen, "Ground-to-aerial image geo-localization with a hard exemplar reweighting triplet loss," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 8391–8400.
- [9] H. Yang, X. Lu, and Y. Zhu, "Cross-view geo-localization with layer-to-layer transformer," *Advances in Neural Information Processing Systems*, vol. 34, pp. 29 009–29 020, 2021.
- [10] D. Yuan, F. Maire, and F. Dayoub, "Cross-attention between satellite and ground views for enhanced fine-grained robot geo-localization," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2024, pp. 1249–1256.
- [11] Z. Xia, O. Booi, M. Manfredi, and J. F. Kooij, "Visual cross-view metric localization with dense uncertainty estimates," in *European Conference on Computer Vision*. Springer, 2022, pp. 90–106.
- [12] S. Zhu, T. Yang, and C. Chen, "Vigor: Cross-view image geo-localization beyond one-to-one retrieval," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 3640–3649.
- [13] X. Zhang, W. Sultani, and S. Wshah, "Cross-view image sequence geo-localization," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 2914–2923.
- [14] Y. Shi, X. Yu, S. Wang, and H. Li, "Cvlnet: Cross-view feature correspondence learning for video-based camera localization," in *Proceedings of the Asian Conference on Computer Vision*, 2022, pp. 652–669.
- [15] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "Netvlad: Cnn architecture for weakly supervised place recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 5297–5307.
- [16] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1125–1134.
- [17] A. Toker, Q. Zhou, M. Maximov, and L. Leal-Taixé, "Coming down to earth: Satellite-to-street view synthesis for geo-localization," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 6488–6497.

- [18] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [19] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2020.
- [20] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *European conference on computer vision*. Springer, 2020, pp. 213–229.
- [21] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [22] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [23] J. L. Elman, "Finding structure in time," *Cognitive science*, vol. 14, no. 2, pp. 179–211, 1990.
- [24] O. Vinyals, S. V. Ravuri, and D. Povey, "Revisiting recurrent neural networks for robust asr," in *2012 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2012, pp. 4085–4088.
- [25] I. Sutskever, J. Martens, and G. E. Hinton, "Generating text with recurrent neural networks," in *Proceedings of the 28th international conference on machine learning (ICML-11)*, 2011, pp. 1017–1024.
- [26] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," *International journal of computer vision*, vol. 115, pp. 211–252, 2015.
- [27] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, 2010, pp. 249–256.
- [28] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2018.